Invertible Neural Network for Data Privacy Mitigation: Synthetic Time-Series Data Generation

Malgorzata Schwab

University of Colorado, Denver CO 80204 USA
malgorzata.schwab@ucdenver.edu

Abstract. This concept paper is inspired by a two-fold concern of preserving data privacy in the face of rapidly increasing role of machine learning across the wide spectrum of domains, as well as providing an insight into data quality behind the models affecting their fairness. It brings to the forefront the importance of intelligent synthetic data generation to support state-of-the-art model training crowdsourcing and data democratization, while protecting privacy of actual datasets, securing intellectual property enabling a competitive edge, and providing the means to understand and manage biases naturally present in live data.

As part of a larger study aimed at Trustworthy AI and data privacy, we explore the applicability of Invertible Neural Networks to the effective and sustainable synthetic data generation particularly important in the healthcare domain. We research the topic of reversibility in deep neural networks and leverage their remarkable data reconstruction capabilities to build a framework for intelligent synthetic data generation, which also fulfills the data quality verification requirements. We apply previous findings and principles regarding variational autoencoders, deep generative maximum-likelihood training and invertibility in neural networks to propose a configurable network architecture enabling bias-free synthetic data generation demonstrated on the ECG dataset.

We build an Invertible Neural Network (INN) Synthetic Data Generator and craft the term INN-SDG, which we present it in the context of a software-as-a-service systems engineering pattern. The universal and modular invertible network architecture proposed in this paper provides a blueprint to efficiently implement an INN-SDG for any data and across various business domain. It separates the concerns of generating the data and then leveraging it to build and deploy the model, putting the respective tasks into the hands of the parties best suited for their execution.

An INN, which is invertible by construction, offers a superb data generation capability where new data adheres to the same probability density distribution as the original input domain. Being a flow-based model, an INN can learn the underlying structure of a given dataset and generate high-quality new data without complex optimization. Theoretically superior to GANs or Variational Autoencoders, flow-based models learn data distribution explicitly, and good density estimation is essential for synthetic data generation and anomaly detection. Even though other approaches exist, an INN-based synthetic data generator presents an attractive alternative versatile to fulfill both data generation and data validation expectations.

Here, a pre-trained Invertible Neural Network is leveraged to generate new data in the reverse flow's invocation based on the augmented randomized sampling from the learned feature distribution in the input domain, resembling latent space. Random samples truncated to match the artificial bottleneck of the trained INN, presented to the inverse flow, turn into a synthetic representation of the training dataset.

The generated sample undergoes validation to determine its quality measured as a resemblance to the original dataset and presented as the sample's Pertinence Score. If it is outside of the acceptable margin configurable for a given use case, the generated sample shall be discarded.

$$|| X_{generated} - X_{original}|| < \text{Pertinence Margin}$$

Data quality validation problem could be approached as anomaly detection and leverage the same INN already trained as autoencoder, rejecting the synthetic samples that do not resemble the real-life dataset, on which the network was originally trained. The proposed INN Synthetic Data Generator is thus readily available to serve as a new sample validator in its regular pass forward inference flow. This novel architecture delivers a possibility of a self-contained black-box synthetic data generation machine learning appliance and provide a generalizable "blue ocean" solution to data privacy preservation.

An INN-SDG fulfills privacy-preserving machine learning (PPML) principles, which encompass safeguarding sensitive data during AI models' training and deployment, ensuring that confidential information remains secure and anonymous throughout machine learning lifecycle. PPML techniques include differential privacy to protect the data points within a dataset by adding noise, or data anonymization paradigm, which entails modifying data to expunge personally identifiable information (PII) and remove any connections to specific individuals, while preserving referential integrity. With the INN-based architecture, new samples are drawn based on the learned probability distribution infused with randomness, so exact resemblance to the actual training data points, although not impossible, would be coincidental. Privacy preserving principles are thus upheld.

Keywords: Invertible, Synthetic, Unbiased, Faithful, Privacy, Trustworthy